



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

TU

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/698,182	10/30/2003	Magnus Karlsson	200311961-1	2501
22879	7590	05/02/2006		
HEWLETT PACKARD COMPANY P O BOX 272400, 3404 E. HARMONY ROAD INTELLECTUAL PROPERTY ADMINISTRATION FORT COLLINS, CO 80527-2400				EXAMINER IWASHKO, LEV
				ART UNIT 2186 PAPER NUMBER

DATE MAILED: 05/02/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/698,182	KARLSSON ET AL.	
	Examiner Lev I. Iwashko	Art Unit 2186	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) Responsive to communication(s) filed on 13 March 2006.
- 2a) This action is FINAL. 2b) This action is non-final.
- 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) Claim(s) 1-27 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) Claim(s) _____ is/are allowed.
- 6) Claim(s) 1-27 is/are rejected.
- 7) Claim(s) _____ is/are objected to.
- 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) The specification is objected to by the Examiner.
- 10) The drawing(s) filed on 30 October 2003 is/are: a) accepted or b) objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 - a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|---|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date <u>02/09/2004</u> | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION***Response to Amendment***

1. The amendments to the Specification have been acknowledged.
2. Claims 1-27 stand rejected.

Claim Rejections - 35 USC § 102

3. The following are quotations of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

4. Claims 1-5, 11-14, and 24-27 are rejected under U.S.C. 102(b) as being anticipated by Ye (US Patent 6,374,227 B1)

- Claim 1. A method (*Title – States that this is a method*)
- of selecting a heuristic class for data placement in a distributed storage system comprising the steps of: (*Abstract, line 1 and 12-14 – State that this optimizes allocation of a resource through the specification of a first heuristic*)
- forming an integer program for each of a plurality of heuristic classes, (*Abstract, lines 4-6 – State that an integer program is received for each of a plurality of bids*)
- each of the integer programs comprising an objective of minimizing a replication cost; (*Abstract, lines 14-17 – State that the integer program optimizes allocation, which means that replication costs must therefore be minimized*)
- solving each of the integer programs (*Abstract, lines 6-7 – State that a solution is generated*)
- which provide the replication cost for each of the heuristic classes; (*Abstract, lines 17-22 – State that there is a maximization problem*)
- and selecting the heuristic class having a low replication cost. (*Abstract, lines 7-10 – State that there is an optimizer engine that is coupled to the file and solver, so in essence, the heuristic class with a low replication cost is selected*)
- Claim 2. A method (*Title – States that this is a method*)

- of selecting a heuristic class for data placement in a distributed storage system comprising the steps of: (*Abstract, line 1 and 12-14 – State that this optimizes allocation of a resource through the specification of a first heuristic*)
 - forming a general integer program which models the data placement; (*Column 21, line 67 and Column 22, line 1 – State that there is a standard integer program according to the data reflecting the optimization problem*)
 - forming a specific integer program which models a heuristic class for the data placement, (*Column 22, lines 42-46 – Declare an enhanced integer program that is generated by optimizer engine 16. Column 23, lines 6-7 – State that optimizer engine 16 applies the heuristic step*)
 - the general and specific integer programs each comprising an objective of minimizing a replication cost; (*Column 17, lines 40-53 – Explains how if there are the same data, then cuts are made*)
 - solving the general integer program which provides a general lower bound for the replication cost; (*Column 7, lines 49-52 – State that there is a lower bound to the integer program solution provided by the LP relaxation solution for the optimization problem*)
 - solving the specific integer program which provides a specific lower bound for the replication cost; (*Column 9, lines 44-49 – Disclose a theoretical lower bound which more closely approximates the optimal solution*)
 - and selecting the heuristic class if a difference between the general lower bound and the specific lower bound is within an allowable amount. (*Column 15, lines 50-64 – State that the first heuristic is applied if there is a valid value between the upper and lower cutoff*)
- Claim 3. The method of claim 2 wherein inputs used in the steps of forming the general and specific integer programs comprise a system configuration, (*Column 6, line 67 and Column 7, lines 1-2 – State that there is a processor and a memory according to the operation of the optimizer*)
- a workload, and (*Column 6, lines 47-51 – State that there is a user that requests info*)
 - a performance requirement. (*Column 21, lines 29-41 – State the idea that there are 2 things that are compared, and if assigned priorities, so there will be a ration of positive to negative assignments*)
- Claim 4. The method of claim 3 wherein the performance requirement comprises a bi-modal performance metric. (*Column 21, lines 29-41 – State the idea that there are 2 things that are compared, and if assigned priorities, so there will be a ration of positive to negative assignments*)
- Claim 5. The method of claim 4 wherein the bi-modal performance metric comprises a criterion and a ratio of successful attempts to total attempts. (*Column 21, lines 29-41 – State the idea that there are 2 things that are*

compared, and if assigned priorities, so there will be a ration of positive to negative assignments)

- Claim 11. The method of claim 3 wherein the general integer program comprises general constraints which model the data placement irrespective of the heuristic class for the data placement. (*Column 5, lines 21-26 – State that there is an integer program that comprises constraints*)
- Claim 12. The method of claim 11 wherein the general constraints comprise a performance constraint which models the performance requirement. (*Column 21, lines 56-58 – State “The prioritizing process may be particularly effective in generating acceptable solutions in situations in which relatively strong carrier selection constraints exist.”*)
- Claim 13. The method of claim 11 wherein the specific integer program comprises the general constraints and a specific constraint. (*Column 5, lines 21-26 – State that there is an integer program that comprises constraints*)
- Claim 14. The method of claim 12 wherein the specific constraint comprises a storage constraint. (*Column 3, lines 36-39 – State that there are capacity constraints*)
- Claim 24. A method (*Title – States that this is a method*)
- of selecting a heuristic class for data placement in a distributed storage system comprising the steps of: (*Abstract, line 1 and 12-14 – State that this optimizes allocation of a resource through the specification of a first heuristic*)
- forming a general integer program which models the data placement; (*Column 21, line 67 and Column 22, line 1 – State that there is a standard integer program according to the data reflecting the optimization problem*)
- forming a plurality of specific integer programs which model a plurality of heuristic classes, (*Column 17, lines 15 – States that there are integer programs, which inevitably model heuristic classes*)
- the general and specific integer programs each comprising an objective of minimizing a replication cost; (*Column 17, lines 40-53 – Explains how if there are the same data, then cuts are made*)
- solving the general integer program which provides a lower bound for the replication cost; (*Column 7, lines 49-52 – State that there is a lower bound to the integer program solution provided by the LP relaxation solution for the optimization problem*)
- solving the specific integer programs which provides the replication cost for each of the heuristic classes; (*Column 9, lines 44-49 – Disclose a theoretical lower bound which more closely approximates the optimal solution*)
- and selecting a particular heuristic class correlated to a low replication cost if a difference between the lower bound and the low replication cost is within an allowable amount. (*Column 15, lines 50-64 – State*)

that the first heuristic is applied if there is a valid value between the upper and lower cutoff)

- Claim 25. A computer readable memory comprising computer code for implementing a method of (*Column 6, line 67 and Column 7, line 1 – State that there is a computer with a memory*)
- selecting a heuristic class for data placement in a distributed storage system, the method of selecting the heuristic class comprising the steps of: (*Abstract, line 1 and 12-14 – State that this optimizes allocation of a resource through the specification of a first heuristic*)
 - forming an integer program for each of a plurality of heuristic classes, (*Abstract, lines 4-6 – State that an integer program is received for each of a plurality of bids*)
 - each of the integer programs comprising an objective of minimizing a replication cost; (*Abstract, lines 14-17 – State that the integer program optimizes allocation, which means that replication costs must therefore be minimized*)
 - solving each of the integer programs (*Abstract, lines 6-7 – State that a solution is generated*)
 - which provide the replication cost for each of the heuristic classes; (*Abstract, lines 17-22 – State that there is a maximization problem*)
 - and selecting the heuristic class having a low replication cost. (*Abstract, lines 7-10 – State that there is an optimizer engine that is coupled to the file and solver, so in essence, the heuristic class with a low replication cost is selected*)
- Claim 26. A computer readable memory comprising computer code for implementing a method of (*Column 6, line 67 and Column 7, line 1 – State that there is a computer with a memory*)
- selecting a heuristic class for data placement in a distributed storage system, the method of selecting the heuristic class comprising the steps of: (*Abstract, line 1 and 12-14 – State that this optimizes allocation of a resource through the specification of a first heuristic*)
 - forming a general integer program which models the data placement; (*Column 21, line 67 and Column 22, line 1 – State that there is a standard integer program according to the data reflecting the optimization problem*)
 - forming a specific integer program which models a heuristic class for the data placement, (*Column 22, lines 42-46 – Declare an enhanced integer program that is generated by optimizer engine 16. Column 23, lines 6-7 – State that optimizer engine 16 applies the heuristic step*)
 - the general and specific integer programs each comprising an objective of minimizing a replication cost; (*Column 17, lines 40-53 – Explains how if there are the same data, then cuts are made*)
 - solving the general integer program which provides a general lower bound for the replication cost; (*Column 7, lines 49-52 – State that*

(there is a lower bound to the integer program solution provided by the LP relaxation solution for the optimization problem)

- solving the specific integer program which provides a specific lower bound for the replication cost; (*Column 9, lines 44-49 – Disclose a theoretical lower bound which more closely approximates the optimal solution*)
- and selecting the heuristic class if a difference between the general lower bound and the specific lower bound is within an allowable amount. (*Column 15, lines 50-64 – State that the first heuristic is applied if there is a valid value between the upper and lower cutoff*)

- Claim 27. A computer readable memory comprising computer code for implementing a method (*Column 6, line 67 and Column 7, line 1 – State that there is a computer with a memory*)
- of selecting a heuristic class for data placement in a distributed storage system, the method of selecting the heuristic class comprising the steps of: (*Abstract, line 1 and 12-14 – State that this optimizes allocation of a resource through the specification of a first heuristic*)
 - forming a general integer program which models the data placement; (*Column 21, line 67 and Column 22, line 1 – State that there is a standard integer program according to the data reflecting the optimization problem*)
 - forming a plurality of specific integer programs which model a plurality of heuristic classes, (*Column 17, lines 15 – States that there are integer programs, which inevitably model heuristic classes*)
 - the general and specific integer programs each comprising an objective of minimizing a replication cost; (*Column 17, lines 40-53 – Explains how if there are the same data, then cuts are made*)
 - solving the general integer program which provides a lower bound for the replication cost; (*Column 7, lines 49-52 – State that there is a lower bound to the integer program solution provided by the LP relaxation solution for the optimization problem*)
 - solving the specific integer programs which provides the replication cost for each of the heuristic classes; (*Column 9, lines 44-49 – Disclose a theoretical lower bound which more closely approximates the optimal solution*)
 - and selecting a particular heuristic class correlated to a low replication cost if a difference between the lower bound and the low replication cost is within an allowable amount. (*Column 15, lines 50-64 – State that the first heuristic is applied if there is a valid value between the upper and lower cutoff*)

Art Unit: 2186

5. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

6. Claim 6 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of Alvarez et al. (US PG Pub 2002/0177989 A1).

Ye teaches the limitations of claims 2 and 3 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of data access latency.

Ye fails to teach claim 6 which states: "The method of claim 3 wherein the performance requirement comprises a data access latency." However, Alvarez declares that "the composite performance metric contains two values (say bandwidth "w" and latency "l")" (Section 0075, lines 5-6). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Alvarez's Method and Apparatus for Predicting Multi-Part Performability so that the method would include a data access latency metric that, upon its monitoring, would aid in maximizing the system's efficiency.

7. Claim 7 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of Alvarez et al. (US PG Pub 2002/0177989 A1).

Ye teaches the limitations of claims 2 and 3 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of data access bandwidth.

Art Unit: 2186

Ye fails to teach claim 7 which states: “The method of claim 3 wherein the performance requirement comprises a data access bandwidth.” However, Alvarez declares that “the composite performance metric contains two values (say bandwidth “w” and latency “l”)” (Section 0075, lines 5-6). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Alvarez’s Method and Apparatus for Predicting Multi-Part Performability so that the method would include a data access bandwidth metric that, upon its monitoring, would aid in maximizing the system’s efficiency.

8. Claim 8 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of Alvarez et al. (US PG Pub 2002/0177989 A1).

Ye teaches the limitations of claims 2 and 3 for the reasons above.

Ye’s invention differs from the claimed invention in that there is no mention of data update time.

Ye fails to teach claim 8 which states: “The method of claim 3 wherein the performance requirement comprises a data update time.” However, Alvarez declares that there is an “average time required to repair or replace a component” (Section 0101, lines 11-12). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Alvarez’s Method and Apparatus for Predicting Multi-Part Performability so that the method would include a data update time that, upon its monitoring, would aid in maximizing the system’s efficiency.

9. Claim 9 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of Alvarez et al. (US PG Pub 2002/0177989 A1).

Ye teaches the limitations of claims 2 and 3 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of average data access latency.

Ye fails to teach claim 9 which states: "The method of claim 3 wherein the performance requirement comprises an average data access latency". However, Alvarez declares that "In addition, it can include descriptions of the workload demands that will be put on the target system (for example, average request rate in IO/s, mean IO request size, average number of requests that are sequential in LU address space, as described hereinbefore)." (Section 0118, lines 6-10). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Alvarez's Method and Apparatus for Predicting Multi-Part Performability so that the method would include an average data access latency metric that, upon its monitoring, would aid in maximizing the system's efficiency.

10. Claim 10 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of Alvarez et al. (US PG Pub 2002/0177989 A1).

Ye teaches the limitations of claims 2 and 3 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of data availability requirement.

Ye fails to teach claim 10 which states: "The method of claim 3 wherein the performance requirement comprises a data availability requirement." However, Alvarez declares that "A(c) is the availability of component c" (Section 0101, lines 7-8). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the

Allocation of a Resource of Ye with Alvarez's Method and Apparatus for Predicting Multi-Part Performability so that the method would include a data availability requirement metric that, upon its monitoring, would aid in maximizing the system's efficiency.

11. Claim 15 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the NPL Document "Adaptive File Allocation in Distributed Computer Systems" (Hereby referred to as Mahmood).

Ye teaches the limitations of claims 2, 3, 11, and 12 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of a replica constraint.

Ye fails to teach claim 15 which states: "The method of claim 12 wherein the specific constraint comprises a replica constraint." However, Mahmood declares that there is a "Replication Constraint" (p. 357, half-way down the page). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Mahmood's method for Adaptive File Allocation so that the method would include a replication constraint that would allow for a fixed number of replicas to enter the heuristic, thereby maximizing the efficiency of the system.

12. Claim 16 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the NPL Document "Adaptive File Allocation in Distributed Computer Systems" (Hereby referred to as Mahmood).

Ye teaches the limitations of claims 2, 3, 11, and 12 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of a channel constraint.

Ye fails to teach claim 16 which states: “The method of claim 12 wherein the specific constraint comprises a routing knowledge constraint and further wherein the routing knowledge constraint models an extent to which a data storage node knows of replicas of data objects stored on other data storage nodes.” However, Mahmood declares that there is a “Channel Constraint” (p. 357, half-way down the page) which performs the same function as Ye’s routing knowledge constraint. Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Mahmood’s method for Adaptive File Allocation so that the method would include a channel constraint that would allow for the data storage node to know how many replicas were stored on other nodes, so that there would be no redundant replication or re-routing, thereby maximizing the efficiency of the system.

13. Claim 17 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the NPL Document “Adaptive File Allocation in Distributed Computer Systems” (Hereby referred to as Mahmood).

Ye teaches the limitations of claims 2, 3, 11, and 12 for the reasons above.

Ye’s invention differs from the claimed invention in that there is no mention of access knowledge constraint.

Ye fails to teach claim 17 which states: “The method of claim 12 wherein the specific constraint comprises an access knowledge constraint and further wherein the access knowledge constraint models an extent to which a data storage knows of access to replicas of data objects by clients accessing other data storage nodes.” However, Mahmood declares that there is a “Processing Constraint” (p. 357, half-way down the page) which performs the same function as

Ye's access knowledge constraint. Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Mahmood's method for Adaptive File Allocation so that the method would include a Processing Constraint that would allow for the data storage to know of access to replicas of data objects by clients accessing other data storage nodes so that redundancy is minimized, thereby maximizing the efficiency of the system.

14. Claim 18 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the NPL Document "Reusable Strategies for Testing Safety-Critical Systems" (Hereby referred to as Poonawala).

Ye teaches the limitations of claims 2, 3, 11, and 12 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of a history constraint.

Ye fails to teach claim 18 which states: "The method of claim 12 wherein the specific constraint comprises an activity history constraint." However, Poonawala declares that there is a "History Constraint" (p. 138, lines 1-4) which performs the same function as Ye's activity history constraint. Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Poonawala's Reusable Strategies for Testing Safety-Critical Systems so that the method would include a history constraint that would make known all of the past activities of nodes so that redundancy is minimized, thereby maximizing the efficiency of the system.

Art Unit: 2186

15. Claim 19 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the NPL Document “Reusable Strategies for Testing Safety-Critical Systems” (Hereby referred to as Poonawala).

Ye teaches the limitations of claims 2, 3, 11, and 12 for the reasons above.

Ye’s invention differs from the claimed invention in that there is no mention of a reactive placement constraint.

Ye fails to teach claim 19 which states: “The method of claim 12 wherein the specific constraint comprises a reactive placement constraint.” However, Poonawala declares that there is an “Input-Boundary Constraint” (p. 141, lines 16-17) which performs the same function as Ye’s reactive placement constraint. Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Poonawala’s Reusable Strategies for Testing Safety-Critical Systems so that the method would include an Input-Boundary Constraint that would reflect if prefetching is included, thereby reducing redundancy and maximizing the efficiency of the system.

16. Claim 20 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the Wong et al. (US PG Pub 2002/0184555).

Ye teaches the limitations of claims 2 and 3 for the reasons above.

Ye’s invention differs from the claimed invention in that there is no mention of coupling nodes to networks.

Ye fails to teach claim 20 which states: “The method of claim 3 wherein the system configuration comprises a plurality of data storage nodes coupled by a plurality of network links.” However, Wong declares a “cluster computer system comprising a plurality of nodes”

(Section 0008, lines 3-4). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Wong's Systems and Methods for Providing Diagnostic Services for a Cluster Computer System so that the method would include an multiple nodes connected to multiple network links in order to minimize network congestion, thereby maximizing the efficiency of the system.

17. Claim 21 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the Wong et al. (US PG Pub 2002/0184555).

Ye teaches the limitations of claims 2, 3 and 20 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of coupling nodes to clients.

Ye fails to teach claim 21 which states: "The method of claim 20 wherein the system configuration further comprises a plurality of clients coupled to the data storage nodes." However, Wong declares that a "plurality of nodes may provide an application to a plurality of clients" (Section 0008, lines 4-5). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Wong's Systems and Methods for Providing Diagnostic Services for a Cluster Computer System so that the method would include an multiple nodes connected to multiple clients in order to minimize network congestion, thereby maximizing the efficiency of the system.

18. Claim 22 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the Wong et al. (US PG Pub 2002/0184555).

Ye teaches the limitations of claims 2, 3, 20, and 21 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of data requesting from clients.

Ye fails to teach claim 22 which states: "The method of claim 21 wherein the workload comprises at least some of the clients requesting data objects stored on the data storage nodes." However, Wong declares that " Each of the plurality of nodes may provide an application to a plurality of clients" (Abstract, lines 5-6). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Wong's Systems and Methods for Providing Diagnostic Services for a Cluster Computer System so that the method would include the clients requesting data objects so that information transfer could take place, thereby making the system functional.

19. Claim 23 is rejected under 35 U.S.C.103(a) as being unpatentable over Ye as applied to claim 3 above, further in view of the Wong et al. (US PG Pub 2002/0184555).

Ye teaches the limitations of claims 2, 3, 20, 21, and 22 for the reasons above.

Ye's invention differs from the claimed invention in that there is no mention of where to store data objects.

Ye fails to teach claim 23 which states: "The method of claim 22 wherein the workload further comprises at least some of the clients storing some of the data objects on the data storage nodes." However, Wong declares that a "The volume management software hide the details about where data is sorted on the drives within nodes" (Section 0065, lines 5-7). Therefore, it would have been obvious to one of ordinary skill in the art to combine the System and Method for Optimizing the Allocation of a Resource of Ye with Wong's Systems and Methods for Providing Diagnostic Services for a Cluster Computer System so that the method would allow

storing objects on nodes in order to make the data quickly accessible by the system, thereby maximizing the overall efficiency of the system.

Response to Arguments

20. Applicant's arguments (filed March 13, 2006) with respect to claims 1-27 have been considered but are moot in view of the previous and following ground(s) of rejection.
21. With regards to Claim 1, the Applicant alleges that "Ye teaches nothing about selecting a heuristic class for data placement in a distributed storage system". The Applicant's Specification states the following: "Each of the heuristic classes represents a range of data placement heuristics. A heuristic comprises a method employed by a computer that uses an approximation technique to attempt to find a near optimal solution but without an assurance that the approximation technique will find a near optimal solution. Heuristics work well at finding the quasi optimum solution provided that a problem definition for a particular problem falls within a range of problem definitions appropriate for a selected heuristic. One skilled in the art will recognize that the term "heuristic" can be employed narrowly to define a search technique that does not provide a result which can be compared to a theoretical best result or it can be employed more broadly to include approximation algorithms which provide a result which can be compared to a theoretical best result. In the context of the present invention, the term "heuristic" is used in the broad sense, which includes the approximation algorithms. Thus, the term "approximation technique" should be read broadly to refer to both heuristics and approximation algorithms" (Page 5, lines 15-29). With the above having been specified, a heuristic class can then simply be a range of data placement techniques that attempt to find a near optimal solution. Ye states the following: "Software for optimizing the allocation of a resource, the software being

embodied in a computer-readable medium and when executed operable to: access resource allocation data comprising a demand for allocation of the resource, a plurality of bids for the resource, and a plurality of reserve bids for the resource; construct a maximization problem according to the demand, the bids, and the reserve bids, the bids being treated as binary variables and the reserve bids being treated as continuous variables, the maximization problem incorporating at least one problem-specific constraint that corresponds to a special ordered set and reflects that only a limited quantity may be awarded based on reserve bids placed in multiple bidding rounds by a particular entity;” (Column 36, lines 48-63). Therefore, the Applicant’s argument is moot in view of the prior art.

22. Further in regards to Claim 1, the Applicant states that Ye does not teach “forming an integer program for each of a plurality of heuristic classes”, “minimizing a replication cost”, “solving each of the integer programs which provide the replication cost for each of the heuristic classes, or “selecting the heuristic class having a low replication cost”. With the first limitation (involving the question of selecting a heuristic class for data placement in a distributed storage system) having been proved in the above response to Applicant’s argument, the preceding arguments are now moot in view of the prior art and the original 102 rejections above fully apply.

23. With regards to Claim 2, the Applicant argues that Ye “does not teach forming an integer program which models the data placement”. However, Ye states the following, which does indeed teach the limitation in the claim: “A solver (18) receives an integer program and generates an LP relaxation solution to the integer program. An optimizer engine (16) coupled to the file (14) and to the solver (18) receives the data and the LP relaxation solution and generates

an enhanced integer program that includes at least one cut according to the data. The cut includes a lifted cover inequality of a specified general form that the LP relaxation solution violates. A specified parameter associated with the lifted cover inequality is determined according to a first heuristic. The solver (18) generates a solution to the enhanced integer program that optimizes the allocation of the resource subject to the demand, bids, and reserve bids" (Abstract, lines 5-17). Therefore, the Applicant's argument is moot in view of the prior art.

24. Further in regards to Claim 2, the Applicant states that "Ye does not teach an objective of minimizing a replication cost". However, there is no explicit explanation of how the replication cost needs to be minimized, or what exactly comprises a replication cost. Therefore, the following rejection still applies, as it describes if there is the same data, then cuts are made: "Further suppose that bids corresponding to $x[1]$ and $x[4]$ are the same bid submitted in different rounds and therefore constitute a special ordered set having the problem-specific constraint $x[1]+x[4].ltoreq.1$. In this case, although the most optimal integral solution is $x[1]=1$ and $y[6]=50$, the LP relaxation solution will instead yield $x[1]=1$, $x[2]=0.25$, and $x[3]=x[4]=x[5]=y[6]=0$. This solution, which includes the fractional value $x[2]=0.25$, is undesirable since it cannot in reality be implemented because $x[2]$ is a binary variable that must have "0" or "1" as its value. Optimizer engine 16 generates at least two cuts in this example--a first cut in the form of a violated cover inequality strengthened according to a special ordered set and a second cut in the form of a violated lifted cover inequality--to prevent IP solver 18 from generating this fractional solution" (Column 17, lines 40-53). Therefore, the Applicant's argument is moot in view of the prior art.

Art Unit: 2186

25. Further in regards to Claim 2, the Applicant states that “Ye does not teach solving an integer program which provides a specific lower bound for the replication cost. However, Ye states “the LP relaxation solution provides a lower bound to the integer program solution that all real solutions to the optimization problem must equal or, as is more likely in typical optimization scenarios, exceed.” (Column 7, lines 49-53). Therefore, the Applicant’s argument is moot in view of the prior art.

26. Further in regards to Claim 2, the Applicant states that “Ye does not teach solving an integer program which provides a general lower bound for the replication cost. However, Ye states “This cutting process also improves the lower bound that the LP relaxation solution provides; that is, cuts introduced according to the cutting process allow the LP relaxation solution for root node 32--the theoretical lower bound on the total shipping cost--to more closely approximate the optimal real solution” (Column 9, lines 44-49). Therefore, the Applicant’s argument is moot in view of the prior art.

27. Further in regards to Claim 2, the Applicant states that “Ye does not teach selecting the heuristic class if a difference between the general lower bound and the specific lower bound is within an allowable limit”. However, even if the Applicant is not satisfied with the previous rejection above, a second reference to Ye is provided as follows: “Optimizer engine compares the unit prices of there bids at step 402 and selects the bid with the lowest unit price at step 404” (Column 23, lines 63-65). Therefore, the Applicant’s argument is moot in view of the prior art.

28. With regards to Claim 24, the Applicant states that “Ye does not teach forming a plurality of specific integer programs which model a plurality of heuristic classes”. As per the above 102 rejections and previous answers to Applicant’s arguments, line 15 of Column 17 in Ye discloses

a plurality of integer programs. Therefore, the Applicant's argument is moot in view of the prior art.

29. Further in regards to Claim 24, the Applicant states that "Ye does not teach "a heuristic class, heuristic classes, or forming a plurality of specific integer programs which model a plurality of heuristic classes". As per the above 102 rejections and previous answers to Applicant's arguments, line 15 of Column 17 in Ye discloses a plurality of integer programs. Therefore, the Applicant's argument is moot in view of the prior art.

30. With regards to Claim 24, the rest of the Applicant's argument are moot in view of the prior art.

31. Claims 3-23 remain rejected as they are directly dependent on Claim 2, and no further arguments were presented to invalidate the rejections.

32. Claims 25-27 remain rejected as they are directly dependent on Claim 1, and no further arguments were presented to invalidate the rejections.

Conclusion

33. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

34. A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event,

however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

35. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Lev I. Iwashko whose telephone number is (571)272-1658. The examiner can normally be reached on (Mon-Fri), from 8-4PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matt Kim can be reached on (571)272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).



Lev Iwashko



MATTHEW KIM
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100